

Work package 2, Deliverable D 203

Allowing Exchange, Integration and Analysis of IWRM Data via the Semantic Web

Benedikt Kämpgen ⁽¹⁾

David Riepl ⁽²⁾

Bernd Herrmann ⁽³⁾

Denny Vrandečić ⁽¹⁾

Andreas Harth ⁽¹⁾

⁽¹⁾ Karlsruhe Institute of Technology (KIT) – Institute of Applied Informatics and Formal Description Methods

⁽²⁾ Karlsruhe Institute of Technology (KIT) – Institute of Applied Geoscience

⁽³⁾ Helmholtz-Centre for Environmental Research – UFZ

April 24th, 2013

Corresponding author:

Benedikt Kämpgen
KIT, AIFB
Englerstr. 11
benedikt.kaempgen@kit.edu
+49 (721) 608 48941

Sponsored by



Federal Ministry
of Education
and Research

Funded by the German Federal Ministry of Education and Research (BMBF):

FKZ 02WM1079

Project Coordination:

Main Coordinator:	Assistant Coordinators	
Prof. Dr. Nico Goldscheider Institute of Applied Geosciences, Karlsruhe Institute of Technology Adenauerring 20b 76131 Karlsruhe Germany Phone: +49 (0) 721 60845465 Fax: +49 (0) 721 606 279	Prof. Dr. Martin Sauter Department of Applied Geology Göttingen University Goldschmidtstrasse 3 37077 Göttingen Germany Phone: +49 (0)551 39 79 11 Fax: +49 (0)551 39 93 79	Dr. Roland Mueller & Dr. Stefan Geyer Helmholtz Centre for Environmental Research (UFZ) Permoserstr. 15 04318 Leipzig Germany Phone: +49 (0)341 235 30 00 Fax: +49 (0)341 235 2885

<http://www.iwrm-smart2.org/>

Contents

Contents	3
1. Introduction	4
2. Wadi Shueib IWRM Scenario.....	6
3. Publishing and Using IWRM Data in the Semantic Web.....	10
1.1 IWRM Ontology.....	10
1.2 DROPEDIA as Linked Data.....	13
1.3 SMART-DB as Linked Data	15
1.4 Integrated SMART Knowledge Base	21
1.5 SMART Data Explorer	25
4. Discussion of Results	31
5. Related Work.....	33
6. Conclusions.....	34
7. Bibliography	35

1. Introduction

Several regions in the world suffer from insufficient natural freshwater resources. A critical region lies at the Lower Jordan River where a steadily increasing population from Israel, Jordan and Palestine has access only to a constantly decreasing amount of physical water. In a collaborative effort by scientists and decision makers from Israel, Jordan, Palestine and Germany, the BMBF-sponsored SMART project conducts research and development on Integrated Water Resources Management (IWRM) against water scarcity in the region of the Lower Jordan Valley. To achieve a Sustainable Management of Available Water Resources in the Lower Jordan Rift Valley (SMART), SMART intends to innovate water resources management – inter alia – through more interoperable information systems.

Different from other water resources management methods, IWRM processes consider social, economical and ecological objectives simultaneously when comparing and assessing long-term planning scenarios in a catchment study area. Common IWRM challenges are 1) a lack of operational guidelines to help scientists contributing to IWRM processes and 2) missing knowledge management methods and tools to share and integrate information from social, economical and ecological sciences. To overcome those challenges in the future, this work contributes to the following goals:

- Create and maintain a knowledge base for IWRM data. Such a knowledge base needs to be partly filled automatically with data from available data sources such as relational databases (e.g., SMART-DB) and partly filled manually by scientists (e.g., DROPIEDIA). Both highly structured data from relational databases and unstructured data from text is to be stored in the knowledge base.
- Define exchange format protocols to 1) allow queries to the knowledge base by third parties, e.g., application developers (e.g., data analysis and visualisations tools) and 2) allow easily extending the knowledge base with data from third parties, e.g., other IWRM-related projects (e.g., GLOWA, EXACT, SUMAR).

To fulfil these goals, we make use of concepts and technologies from Semantic Web research and applications. The Semantic Web (Berners-Lee, Hendler, & Lassila, 2001) is an effort to extend the current Web so that contained information can be understood by machines and processed automatically for the benefit of human users, e.g., when searching for information and collaborating with others.

More concretely, this work contains the following contributions:

- We formalise the IWRM knowledge and decision support domain using an OWL ontology. OWL (Group, 2012) is based on RDF (Manola & Miller, 2004), a flexible, graph-based data exchange format, and is a standard language for representing, sharing and automatic processing information on the Semantic Web.
- We propose to use the Linked Data principles (Berners-Lee, Linked Data, 2006) as a suitable exchange format protocol for publishing and consuming IWRM data and provide both SMART-DB and semi-structured data from DROPEdia as Linked Data reusing our IWRM ontology.
- We present an ETL pipeline to automatically create and regularly update an IWRM knowledge base in a graph-based database.
- We design and develop the SMART Data Explorer as an application for visualising and analysing integrated IWRM data.

Our work is primarily related to the following SMART deliverables:

- For an automatic comparison of scenarios (alternatives) according to indicators (criteria), our IWRM process relies on the SMART decision support system (DSS) of Multi Criteria Decision Making (MCDM) tools (Bensabat, Kitron-Belinkov, Milgrom, & Dror, 2012), (Bensabat & Dror, Multi-Criteria Optimization (MCO) Module - Version 1.2, 2013).
- The Knowledge Management System, DROPEdia (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013), is published as Linked Data and integrated in the IWRM knowledge base.
- Data from the SMART database (SMART-DB) with Data and Information System (DAISY) (Geyer & Siebert, 2009) is published as Linked Data and integrated in the IWRM knowledge base.

The remainder of this work is structured as follows: In Section 2, we describe a concrete IWRM process that we use as a motivational and exemplary scenario throughout the work. In Section 3, we describe our approach comprising of 1) an IWRM ontology, 2) mechanisms to publish DROPEdia and the SMART-DB and as Linked Data reusing the IWRM ontology, 3) an ETL pipeline to create and maintain an IWRM knowledge base with integrated SMART data and 4) the SMART Data Explorer to analyse data from the IWRM knowledge base. In Section 4, we discuss the solution to our scenario. In Section 5, we describe related work, after which, in Section 6, we summarise the work.

2. Wadi Shueib IWRM Scenario

Riepl (Riepl, 2013) has presented a use case of the development of an IWRM decision process for the Wadi Shueib catchment in Jordan. The Wadi Shueib IWRM Decision Process we first define as a well-known decision support problem and second we describe the information that is needed to solve this problem with common decision support tools.

An IWRM decision problem is characterised by a large set of alternative Water Management strategies with multiple, partly conflicting evaluation criteria. As such, the problem of finding consensus and a final choice can be described as a multi-criteria-decision-analysis problem (Bensabat, Kitron-Belinkov, Milgrom, & Dror, 2012). The typical approach for a multi-criteria-decision-analysis problem (MCDA) is for the decision maker(s) to assign weights/preferences to criteria and to afterwards rank the decision alternatives with respect to their performance according to criteria and weights. Bensabat et al. (Bensabat, Kitron-Belinkov, Milgrom, & Dror, 2012) (Bensabat & Dror, Multi-Criteria Optimization (MCO) Module - Version 1.2, 2013) have presented algorithms and tools for MCDA, e.g., for the Analytical Hierarchy Process (AHP).

Now, we describe the Wadi Shueib IWRM Decision Process (Wadi Shueib Process) in more detail. In particular, we describe some of the decision alternatives (scenarios) and criteria (indicators) according to which the alternatives later get ranked. In our use case, decision makers instantiate the Wadi Shueib IWRM Decision Process based on the national water strategy of Jordan. They further define IWRM objectives that are to be optimised in the process, e.g., to “Increase volume of captured and treated wastewater”.

Domain experts from social, economical and ecological sciences define indicators (criteria in MCDM) to evaluate the suitability of scenarios according to IWRM objectives. For instance, to evaluate the “Increase volume of captured and treated wastewater” of a strategy, the “Municipal Waste Water Treatment Ratio” is suitable. The indicator quantifies the potential impact of discharged untreated waste water in terms of health and resources protection. Also, domain experts define scenarios (decision alternatives in MCDA), a realistic future state of the catchment. An example scenario is "Wadi Shueib BAU" the water strategy implementation according to the current plans of the Jordanian national water strategy. It is calculated from an average of two more detailed scenarios “Wadi Shueib HRP-BAU” assuming a continuous dry period, high population growth, and strong demand development, and “Wadi Shueib LRP-BAU” assuming no dry period and low resources pressure (LRP) according to population growth and water demand.

The process of defining indicators and scenarios is a highly collaborative and knowledge intensive process: social, economical and ecological information needs to be considered simultaneously and strong consensus found.

Domain experts contribute expert analyses such as a data-preprocessing analysis “Comparison of areal rainfall interpolations and literature values for Wadi Shueib area”, estimations and assumptions on “Population growth for the Wadi Shueib area 2010-2025 with growth rates are according to the HPC-Scenarios I and III”, and the evaluation of indicators for scenarios in “Municipal Waste Water Treatment Ratio in Wadi Shueib Scenarios”. Domain experts are motivated to contribute their analyses 1) if analyses are published for the entire IWRM community for citations, feedback and possible future collaborations 2) if analyses can be compared and integrated with other analyses to generate additional knowledge by more easily combine information from different sources and 3) if applications can be developed on top of published information for innovative usages.

In particular, some expert analyses require information from DROPEDIA and SMART-DB. However, DROPEDIA (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013) has been created by and is available from a different institution than the SMART-DB (Geyer & Siebert, 2009). As such, access mechanisms, data format and data schema differ, e.g., access to the SMART-DB is restricted to certain user-interactions in DAISY and the data in the SMART-DB is stored in a relational data format on a database based on Oracle SQL server. Furthermore, although DROPEDIA and SMART-DB share the concept “location”, a direct link between individual locations is not directly possible, since DROPEDIA and SMART-DB use different identifiers, e.g., DROPEDIA uses “Baqqouria Spring” and SMART-DB uses “AM0530”. Thus, for comparing information from both platforms, domain experts often need additional efforts such as manually copy-pasting of tables and additional tools such as Microsoft Excel.

Also, decision makers use the WEAP software¹ to simulate scenarios and need to document WEAP models (e.g., as maps). Climate time series from DROPEDIA or SMART-DB are interesting to be imported into WEAP models.

Finally, the decision makers fill a decision matrix with the information provided by the domain experts for input in a multi-criteria-decision-analysis tool (Bensabat, Kitron-Belinkov, Milgrom, & Dror, 2012). For instance, Figure 1 illustrates the decision matrix for the Wadi Shueib Process comparing scenarios “Business As Usual” (BAU), “Full Implementation” (FI) and “Reference”

¹ <http://www.weap21.org/>

(Ref) according to normalised calculations of indicators “Waste Water Recharge Ratio” (WWrecharge), “Available Groundwater” (AWRinternal), “Municipal Shortage” (WSSmunicipal), “Municipal Supply Requirement” (WSRmunicipal), “Unit Cost” (FSCuc/municipal) and “Environmental Water Stress” (WSIrf). Depending on the weights for indicators, a MCDA tool would rank the scenarios as a possible solution to this Wadi Shueib IWRM decision process.




Figure 1 Decision matrix of Wadi Shueib Process visualised in a line diagram (Riepl, Knowledge-Based Decision Support for Integrated Water Resources Management with an application for Wadi Shueib, Jordan, 2013).

In the following, we derive requirements of the Wadi Shueib Process for an Exchange, Integration and Analysis concept for IWRM data.

Requirements

- It should be possible to define IWRM decision support elements such as IWRM processes, objectives, indicators and scenarios.
- It should be possible to profoundly describe IWRM decision support elements via links to unambiguously defined IWRM knowledge objects such as catchments (e.g., Wadi Shueib), municipalities (e.g., As-Salt) springs (e.g., Baqqouria Spring) and meteorological stations (e.g., Ira Meteo).

- Decision makers and domain experts should be able to collaboratively work (e.g., discuss, comment, re-use) on IWRM decision support elements and IWRM knowledge objects.
- Domain experts should be able to collaboratively work (e.g., discuss, comment, reuse) on expert analyses such as data-preprocessing analyses, estimations and assumption analyses, evaluations of indicators for scenarios. Also, the actual decision matrix to be put in a multi-criteria-decision-analysis tool should be possible to be created by a decision maker.
- All provided information should be published and accessible for the entire IWRM community for citations, feedback and possible future collaborations.
- It should be possible to more easily combine information from different sources to generate additional knowledge, e.g., when comparing different expert analyses. For instance, DROPIEDIA and SMART-DB are two separately created information sources that are to be integrated. Also, WEAP information in the form of maps or timeseries should be possible to be integrated with DROPIEDIA or SMART-DB data as well as calculations (observations) needed for simulations in WEAP queried from DROPIEDIA or SMART-DB. More concrete requirements:
 - A DROPIEDIA webpage that describes an element also contained in SMART-DB shall give a tabular overview of the data that the SMART-DB stores about this element.
 - A DROPIEDIA webpage that describes an element also contained in SMART-DB shall give graphical insights of data available in the SMART-DB.
 - DROPIEDIA should allow to document and discuss datasets and give meta information such as the data owner, a time series visualisation, and an interactive visualisation.
- Not only humans should be able to reuse generated data, but also applications should be possible to develop that automatically process generated data and provide innovative analysis capabilities.

3. Publishing and Using IWRM Data in the Semantic Web

In this section, we describe our approach comprising of 1) an IWRM ontology, 2) mechanisms to publish DROPEdia and the SMART-DB and as Linked Data reusing the IWRM ontology, 3) an ETL pipeline to create and maintain an IWRM knowledge base with integrated SMART data and 4) the SMART Data Explorer to analyse data from the IWRM knowledge base.

1.1 IWRM Ontology

For an IWRM ontology, we have chosen the OWL ontology (Working Group, 2012) for the underlying semantics and RDF (Manola & Miller, 2004) as a data representation format. See Figure 2 for an illustration of the ontology as an RDF graph.

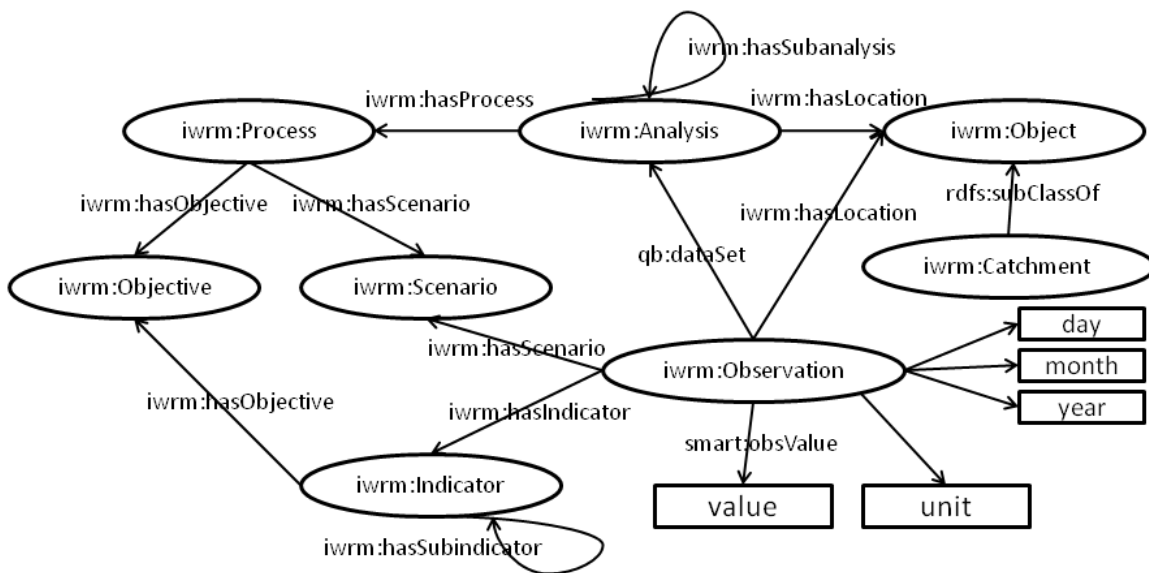


Figure 2 IWRM Ontology

Note, the illustration does only contain the most important concepts and properties of the IWRM ontology; the entire ontology (also called DROPEdia ontology) is more comprehensive and can be browsed in DROPEdia (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013).

An ontology formalised via OWL and RDF brings the following benefits:

- OWL is based on Description Logics (Hitzler, Krötzsch, & Rudolph, 2010) and as such has formal semantics with which computer can derive implicit knowledge. For instance,

from “iwrn:Catchment rdfs:subClassOf iwrn:Object” a machine can automatically derive that any individual of type iwrn:Catchment (e.g., Wadi Shueib) also is of type iwrn:Object. Using “dropedia:Baqqouria_Spring owl:sameAs smart-db:AM0530” a machine “knows” that any property that is valid for Baqqouria Spring in DROPEdia also is valid for the individual represented by AM530 in the SMART-DB.

- There already is a large number of ontologies available. If reused, data integration automatically is made easier, since existing applications may understand reused ontologies. For instance, the IWRM ontology is reusing the RDF Data Cube Vocabulary (Cyganiak & Reynolds, 2013) to represent numeric observations (iwrn:Observation) such as sensor data from the SMART-DB.
- Given the URIs for the abbreviating namespaces such as “iwrn” and “owl”, as described in Table 2, every concept, individual and property of the IWRM ontology is published in the Semantic Web and can be shared, linked to, and extended by other organisations. For that, we use the Linked Data principles (Berners-Lee, Linked Data, 2006). For any given concept, instance and property, in Linked Data, we distinguish three URIs,: 1) Resource, the unique identifier of the “thing” in the Semantic Web 2) HTML, a human readable representation of the “thing” in the common Web; can be nicely-displayed by a common Web browser 3) RDF, a machine readable representation of the “thing”; provides structured information that can be understood by machines (alias computers). See Table 1 for an overview. All three URIs can be opened with a common Web browser. The Resource URI does “content negotiation”, i.e., it forwards “users” to the right representation, e.g., a common Web browser is forwarded to the HTML representation, a Semantic Web application is forwarded to the RDF representation.

Table 1 Three URIs per “thing”: 1) resource itself 2) HTML 3) RDF

Description	URI
Resource	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIResolver/Catchment
HTML	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Catchment
RDF	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:ExportRDF/Catchment

- The ontology defines the concepts and properties which can be populated with individuals (instances), e.g., the “Wadi Shueib” instance of an `iwrn:Catchment`. We will later show how to populate this ontology via DROPIEDIA and SMART-DB. The populated IWRM ontology provides the following benefits:
 - First, the links between individuals allows for follow-your-nose browsing similar to common Web browsing. For instance, from the analysis instance of “Population growth for the Wadi Shueib area 2010-2025 with growth rates are according to the HPC-Scenarios I and III”, the user can browse via “`iwrn:hasLocation`” to the analysed instance of “Wadi Shueib”.
 - Second, the ontology provides the schema for querying the data using SPARQL. SPARQL (Working Group W. S., 2013) is a language for querying and manipulating RDF graph content on the Web or in an RDF store.

Table 2 Namespaces used for IWRM ontology and IWRM knowledge base

Namespace	URI
iwrn	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIResolver/
dropedia	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIResolver/
smart	http://smartdbwrap.appspot.com/
smartdb	http://smartdbwrap.appspot.com/
qb	http://purl.org/linked-data/cube#
owl	http://www.w3.org/2002/07/owl
rdfs	http://www.w3.org/2000/01/rdf-schema

The IWRM ontology has been developed and is provided on DROPIEDIA (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013), the collaborative knowledge management system of SMART. The ontology itself only contains the schema information, i.e., the concepts and properties that can be used for actual data. For instance, the ontology provides the concept “`iwrn:Indicator`” and the property “`iwrn:hasObjective`” however does not provide a concrete indicator that has a specific objective. In the next two sections, we describe

how DROPEDIA and SMART-DB fulfil the need to populate the ontology with such instances and how to also publish those instances as Linked Data.

1.2 DROPEDIA as Linked Data

In this section, we describe how DROPEDIA² is used to populate the IWRM ontology with instance data (individuals) such as the iwrn:Catchment “Wadi Shueib”.

DROPEDIA is based on Semantic MediaWiki (Völkel, Krötzsch, Haller, & Studer, 2006) an Open Source semantic wiki software³ based on MediaWiki. Using specific forms for concepts, users of DROPEDIA can create and modify individuals such as “Wadi Shueib”. Structured information can be queried directly within DROPEDIA. For more information, see D606 (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013).

DROPEDIA allows uploading and displaying KML files as well as linking from the map to specific wiki pages for browsing to background information as illustrated in Figure 3.

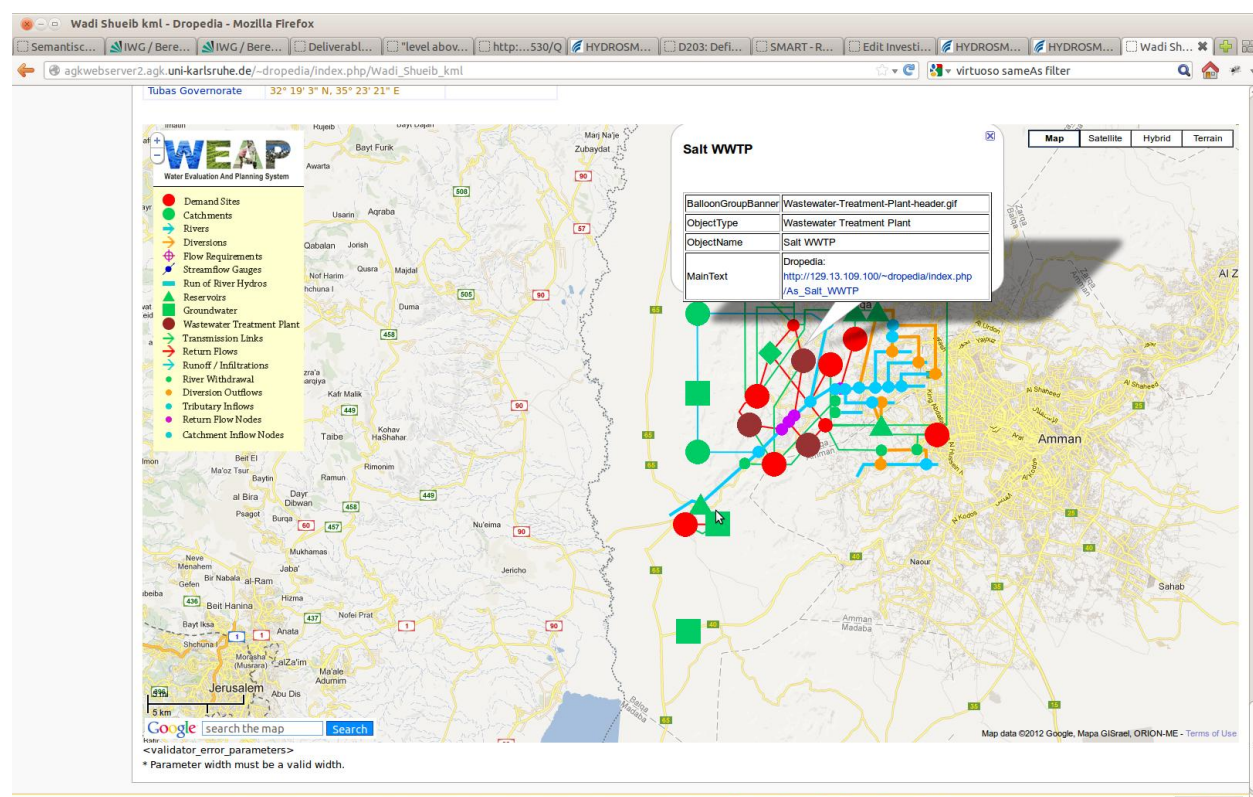
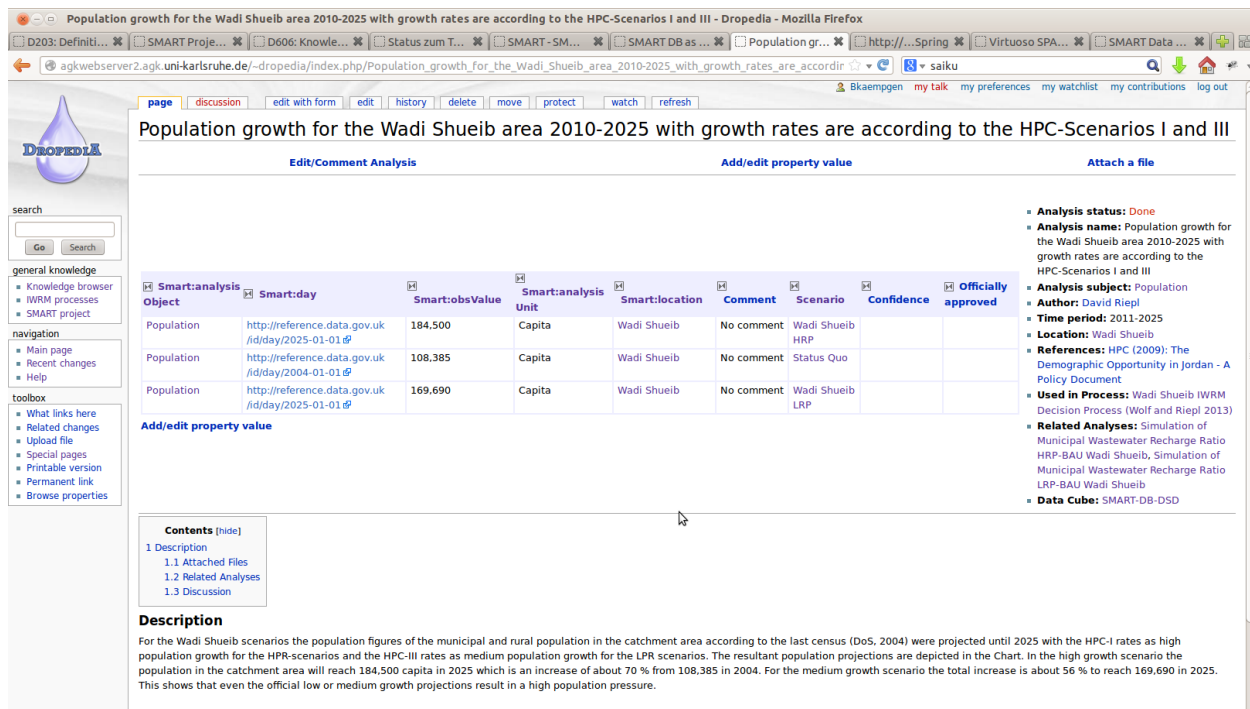


Figure 3 Integration of WEAP KML files with DROPEDIA

² <http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/>

³ <http://www.semantic-mediawiki.org/>

DROPEDIA allows adding observations to expert analyses. For instance, see in Figure 4 created observations for the analysis “Population growth for the Wadi Shueib area 2010-2025 with growth rates are according to the HPC-Scenarios I and III”, where an observation has been created that estimates the population (analysis object/indicator) of Wadi Shueib (location) for the scenario “Wadi Shueib HRP” in 2012 (date) to be 184,500.



Population growth for the Wadi Shueib area 2010-2025 with growth rates are according to the HPC-Scenarios I and III

Smart:analysis Object	Smart:day	Smart:obsValue	Smart:analysis Unit	Smart:location	Comment	Scenario	Confidence	Officially approved
Population	http://reference.data.gov.uk/id/day/2012-01-01	184,500	Capita	Wadi Shueib	No comment	Wadi Shueib HRP		
Population	http://reference.data.gov.uk/id/day/2014-01-01	108,385	Capita	Wadi Shueib	No comment	Status Quo		
Population	http://reference.data.gov.uk/id/day/2025-01-01	169,690	Capita	Wadi Shueib	No comment	Wadi Shueib LRP		

Analysis status: Done

Analysis name: Population growth for the Wadi Shueib area 2010-2025 with growth rates are according to the HPC-Scenarios I and III

Analysis subject: Population

Author: David Riepl

Time period: 2011-2025

Location: Wadi Shueib

References: HPC (2009): The Demographic Opportunity in Jordan - A Policy Document

Used in Process: Wadi Shueib IWRM Decision Process (Wolf and Riepl 2013)

Related Analyses: Simulation of Municipal Wastewater Recharge Ratio HRP-BAU Wadi Shueib, Simulation of Municipal Wastewater Recharge Ratio LRP-BAU Wadi Shueib

Data Cube: SMART-DB-DSO

Contents [hide]

- 1 Description
- 1.1 Attached Files
- 1.2 Related Analyses
- 1.3 Discussion

Description

For the Wadi Shueib scenarios the population figures of the municipal and rural population in the catchment area according to the last census (DoS, 2004) were projected until 2025 with the HPC-I rates as high population growth for the HRP-scenarios and the HPC-III rates as medium population growth for the LRP scenarios. The resultant population projections are depicted in the Chart. In the high growth scenario the population in the catchment area will reach 184,500 capita in 2025 which is an increase of about 70 % from 108,385 in 2004. For the medium growth scenario the total increase is about 56 % to reach 169,690 in 2025. This shows that even the official low or medium growth projections result in a high population pressure.

Figure 4 Example input of observations in DROPEDIA

DROPEDIA covers IWRM knowledge objects such as Catchments, Springs and Wells via a “Knowledge browser”⁴ as well as IWRM decision support elements such as processes, indicators and scenarios via “IWRM processes”.⁵

Just as for the IWRM ontology, developed and provided on DROPEDIA, we publish all structured information as Linked Data and distinguish three URIs per “thing” such as “Wadi Shueib”, see Table 3.

⁴ http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Water_resources_characterisation

⁵ http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/IWRM_processes

Table 3 Three URIs per “thing”: 1) resource itself 2) HTML 3) RDF

Description	URI
Resource	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIResolver/Wadi_Shueib
HTML	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Wadi_Shueib
RDF	http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:ExportRDF/Wadi_Shueib

In Summary, all structured information from DROPEdia is also published as Linked Data. However, it is our goal to not only have access to DROPEdia data, but also to integrate it with SMART-DB data. Then, SMART-DB data could be displayed at appropriate places in DROPEdia and exported for WEAP simulations. Therefore, in the next section, we describe how data from SMART-DB can be exchanged and integrated with other data sources.

1.3 SMART-DB as Linked Data

In this section, we describe how we publish SMART-DB data as Linked Data for reuse. Figure 5 illustrates the architecture of sharing and using SMART-DB data in the Semantic Web.

Using the edit functionality and SPARQL-based (SPARK) query syntax of the DROPEdia, users that understand the SMART ontology (explained in section “SMARTDBWRAP”) can directly create queries to the SMART-DB. Locations in DROPEdia are uniquely identified in SMART-DB using a manually inserted “ID-Code” based on MWI, Mekorot, or PHG. Any user, also without technical background, can visit wiki pages triggering those queries and get visualised the results embedded in the Web pages. Queries are issued to a database that is filled with SMART-DB data. For instance, Qcrumb.com serves a suitable database. Qcrumb.com automatically retrieves the necessary data from the SMART-DB by issuing HTTP requests on a Linked Data Wrapper (SMARTDBWRAP) of the SMART-DB. The SMARTDBWRAP is a Google-App-Engine-based Web application that in turn issues HTTP requests to an XML/HTTP interface (HYDROSMART) on top of the SMART-DB. HYDROSMART translates HTTP requests on predefined URLs into SQL queries to the Oracle relational database and returns the results to SMARTDBWRAP as an XML file describing the results with a fixed schema.

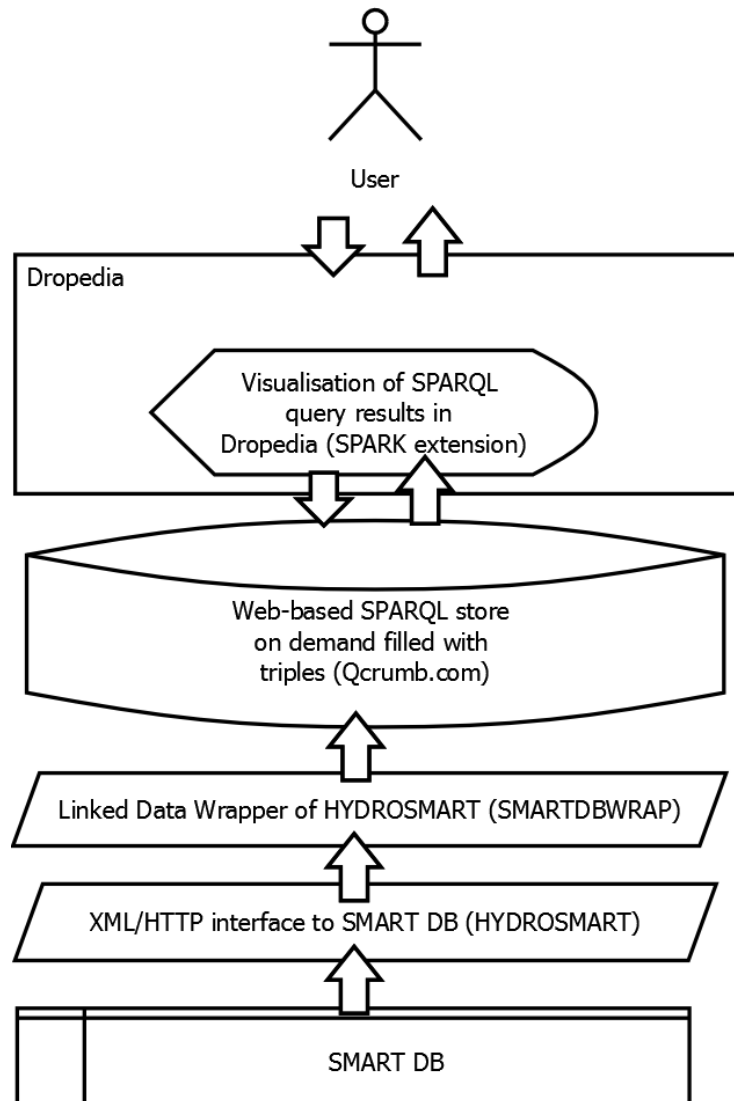


Figure 5 Architecture of publishing and consuming SMART-DB as Linked Data

In the following, we describe in more detail the main parts of the architecture:

SPARK extension

The SPARK extension⁶ allows wiki pages to issue SPARQL queries to any triple store and to display the results in tables and diagrams. SPARK is based on JavaScript. SPARK also allows MediaWiki template parameters such as “`{{FULLPAGENAME}}`” for querying the current wiki page name. This allows to define queries to the triple store that are parameterized by specific information that is provided on a wiki page, e.g., the name of the object represented or a ID code that is used in SMART-DB.

⁶ <http://www.mediawiki.org/wiki/Extension:Spark>

Now, for instance, we can give overviews of SMART-DB data within DROPEDIA such as “Number of analyses” as illustrated in Figure 6; we allow static analyses of SMART-DB data, e.g., comparison of time series as illustrated in Figure 7; and we allow interactive analyses using pivot tables, e.g., comparison of time series as illustrated in Figure 8.

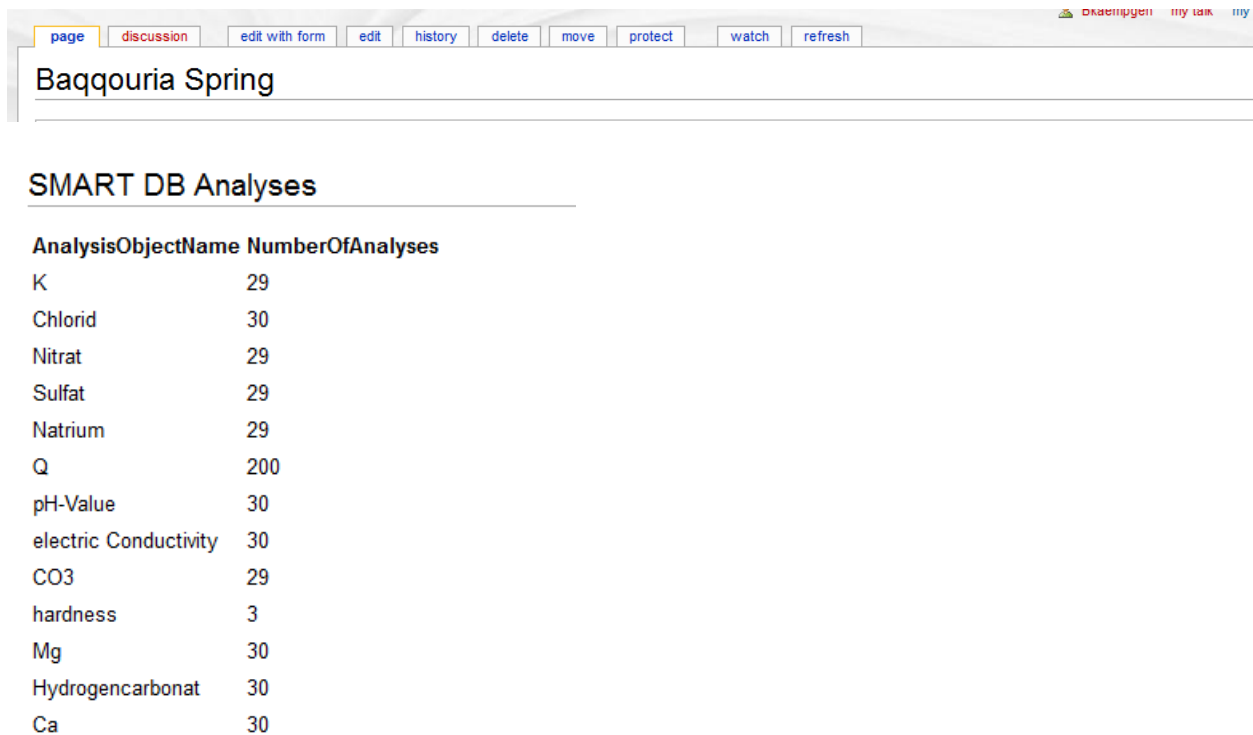


Figure 6 Number of analyses of Baqqouria Spring in SMART-DB




Figure 7 Comparing level above msl at three locations with line chart

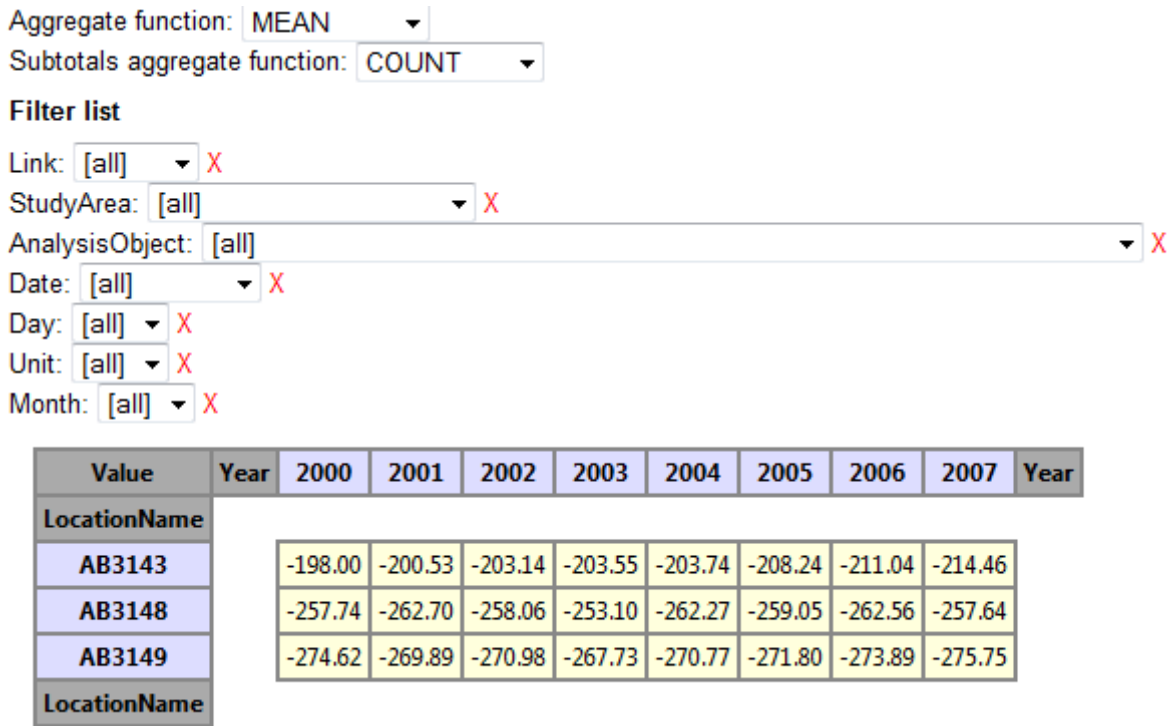


Figure 8 Comparing level above msl at three locations with interactive pivot table

SMARTDBWRAP

SMARTDBWRAP⁷ is based on a Google-App-Engine. It on-the-fly translates XML from an Web interface to SMART-DB (HYDROSMART) into an RDF representation using the IWRM ontology. SMARTDBWRAP provides the following instances. Note, we only mention the resource URI since representation URIs are found via content negotiation:

Description	URI
Dataset of locations	http://smartdbwrap.appspot.com/id/location/ds
Dataset of iwrn:Indicators	http://smartdbwrap.appspot.com/id/analysisobject/ds
Datasets of locations , e.g., "AM0530" ("Baqqouria Spring" in DROPIEDIA)	http://smartdbwrap.appspot.com/id/location/AM0530
Dataset of a location, e.g., "Q" ("Mean Discharge in DROPIEDIA) for "AM0530"	http://smartdbwrap.appspot.com/id/locationdataset/AM0530/Q

In particular, in datasets of a specific location, SMARTDBWRAP publishes instances of iwrn:Observation, i.e., recorded historic sensor data, such as illustrated in Figure 9.

⁷ <http://smartdbwrap.appspot.com/>

```

- <qb:DataSet rdf:about="/id/locationdataset/AM0530/Q">
  <rdfs:label>Locationdata Q@AM0530 Dataset</rdfs:label>
  <qb:structure rdf:resource="http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIRe
  <qb:dataSet rdf:resource="/id/location/AM0530"/>
</qb:DataSet>
- <qb:Observation>
  - <smart:project_Namedataowner>
    Water Information System Information System of Ministry of Water and Irrigation, Jordan. IWRM SMART: Lc
  </smart:project_Namedataowner>
  <smart:study_Area>Wadi Shu eib Basin</smart:study_Area>
  <smart:location rdf:resource="/id/location/AM0530"/>
  <smart:longitude>35,73172</smart:longitude>
  <smart:latitude>31,98404</smart:latitude>
  <smart:day rdf:resource="http://reference.data.gov.uk/id/day/1963-12-01"/>
  <smart:month rdf:resource="http://reference.data.gov.uk/id/month/1963-12"/>
  <smart:year rdf:resource="http://reference.data.gov.uk/id/year/1963"/>
  <smart:longdate>1963-12-01</smart:longdate>
  <smart:date>01.12.63</smart:date>
  <smart:analysis_Object rdf:resource="/id/analysisobject/Q"/>
  <sdmx-measure:obsValue>792</sdmx-measure:obsValue>
  <smart:obsValue rdf:datatype="http://www.w3.org/2001/XMLSchema#decimal">792</smart:obsValue>
  <smart:analysis_Unit>m^3/h</smart:analysis_Unit>
  <qb:dataSet rdf:resource="/id/locationdataset/AM0530/Q"/>
</qb:Observation>

```

Figure 9 Example observation from SMART-DB as RDF/XML

Here, a dataset "Locationdata Q@AM0530 Dataset" publishes an observation from 1963-12-01, at location AM0530 of analysis object (indicator) Q. Note, an observation not simply uses primitive datatypes such as string or date for describing the date, the location, and the analysis object, but instead links to their resources. This assures that machines can disambiguate and find further information about those metadata, e.g., the human readable name of a location.

HYDROSMART

HYDROSMART has been developed by UFZ. It is an HTTP-based interface for structured data from SMART-DB and DAISY (Geyer & Siebert, 2009). SMART-DB is based on an Oracle RDBMS. The goal of the SMART-DB is to store and give access to sensor data, called "locationdata". Figure 10 shows the schema part of SMART-DB that is given access to via HYDROSMART.

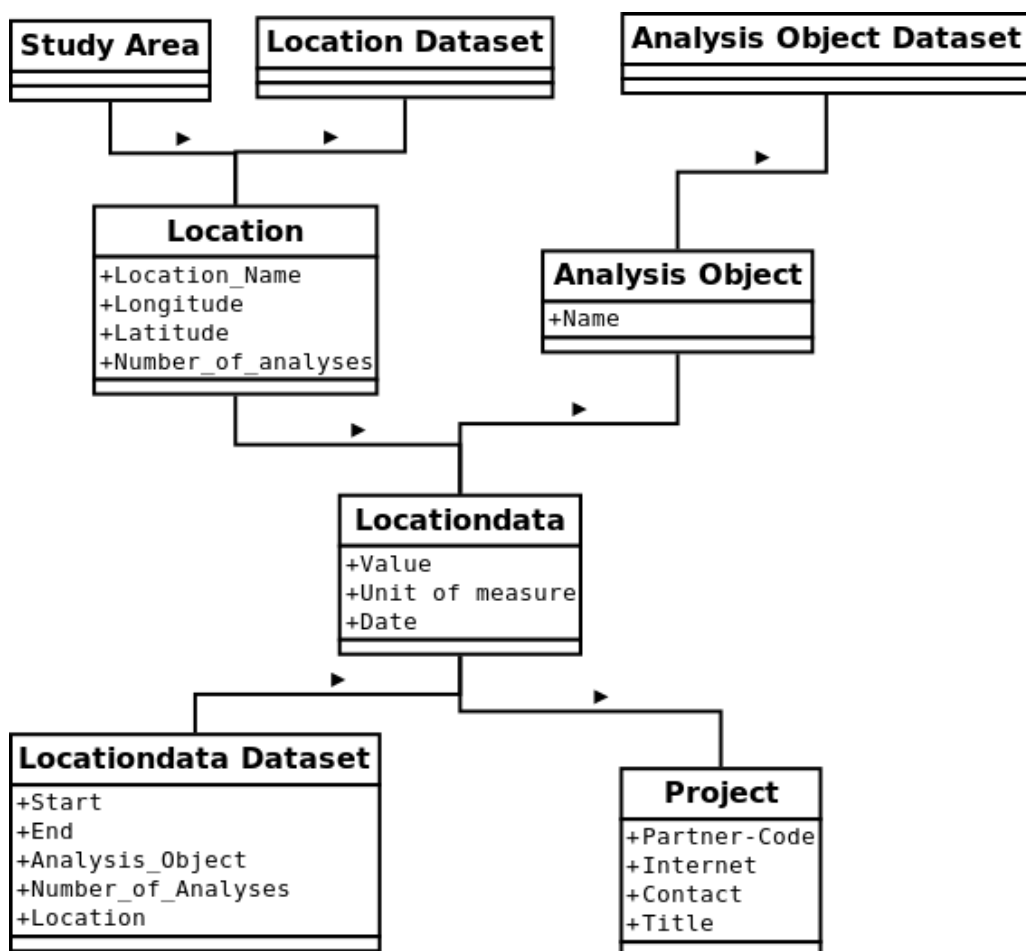


Figure 10 HYDROSMART relevant schema part of SMART-DB

Any locationdata is contained in a locationdata dataset, has been created by a specific project, has been recorded at a specific location, e.g., “AM0530” (“Baqqouria Spring” in DROPEDIA), and measures a certain analysis object, e.g., “Q” (“Mean Discharge” in DROPEDIA). HYDROSMART provides URIs that when resolved issue queries on SMART-DB and return results in a simple XML format. Note the direct correspondence between URIs provided by HYDROSMART and URIs provided by SMARTDBWRAP:

Description	URI
Dataset of locations	http://www2.ufz.de/smarthydro/smartquery?location_name=all
Dataset of iworm:Indicators	http://www2.ufz.de/smarthydro/smartquery?analysis_object_name=all
Datasets of locations , e.g., “AM0530”	http://www2.ufz.de/smarthydro/smartquery?location_data=AM0530
Dataset of a location, e.g., “Q” for “AM0530”	http://www2.ufz.de/smarthydro/smartquery?location=AM0530&analysis_object=Q

Now that SMART-DB data is published as Linked Data, we are able to visualise the SMART-DB data next to DROPEDIA content. Two challenges remain:

1. SMART-DB and DROPEDIA use different identifiers for the same “things”, e.g., AM0530 vs. Baqqouria Spring and Q vs. Mean Discharge. Thus, querying SMART-DB data and DROPEDIA data as a fully integrated data source is not possible.
2. An export of SMART-DB or DROPEDIA data in a WEAP-importable format such as CSV is not directly possible.

In the next two sections we will first present the SMART Integrated Knowledge Base that fully integrates SMART-DB and DROPEDIA data, second, present an analysis application – SMART Data Explorer – that allows CSV exports of integrated data.

1.4 Integrated SMART Knowledge Base

So far, instance data for the IWRM ontology is distributed and is not available from one source. More concretely, SMART-DB data can be visualised within DROPEDIA, yet, querying over the two sources simultaneously is not possible. In particular, SMART-DB and DROPEDIA use different identifiers for identical “things”, e.g., AM0530 for Baqqouria Spring and Q for Mean Discharge. Fully integrating both data sources would mean to allow queries on both data sources simultaneously considering identical elements, i.e., creating a canonical identifier for each group of identical elements. See Figure 11 for an overview of the architecture of an integrated knowledge base for the SMART project.

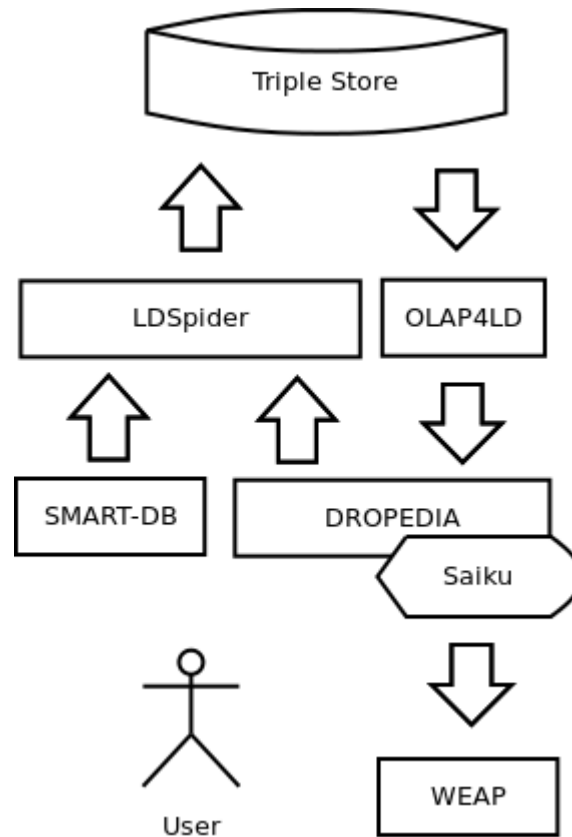


Figure 11 Architecture of Integrated IWRM Knowledge Base

Here, a data warehouse for RDF – a **triple store** (Open Virtuoso)⁸ – is to be automatically and regularly filled with up-to-date data from SMART-relevant data sources. We use **LDSpider**⁹ for automatically “crawling” Linked Data. LDSpider starts with a seed from DROPIEDIA¹⁰ that lists all locations in DROPIEDIA. From a list provided by UFZ, we were able to link identifiers of locations in DROPIEDIA with identifiers of locations in SMART-DB. Therefore, LDSpider can automatically find the respective data about a location in DROPIEDIA also in SMART-DB, as illustrated in Figure 12. Crawled data is then inserted in the triple store. Starting the seed data from DROPIEDIA assures that only those data from SMART-DB is temporarily crawled that is relevant for the Knowledge Management System (Riepl & Kämpgen, Knowledge Management System - Deliverable 606, 2013). This version of relevant data is updated every time the data warehouse (triple store) is populated. It is not the goal to permanently duplicate information from SMART-DB and DAISY, but rather to provide alternative and Data-Integration-focusing access mechanisms that make data owners and generators transparent.

⁸ <http://www.openlinksw.com/wiki/main/>

⁹ <http://code.google.com/p/ldspider/>

¹⁰ <http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Seed>

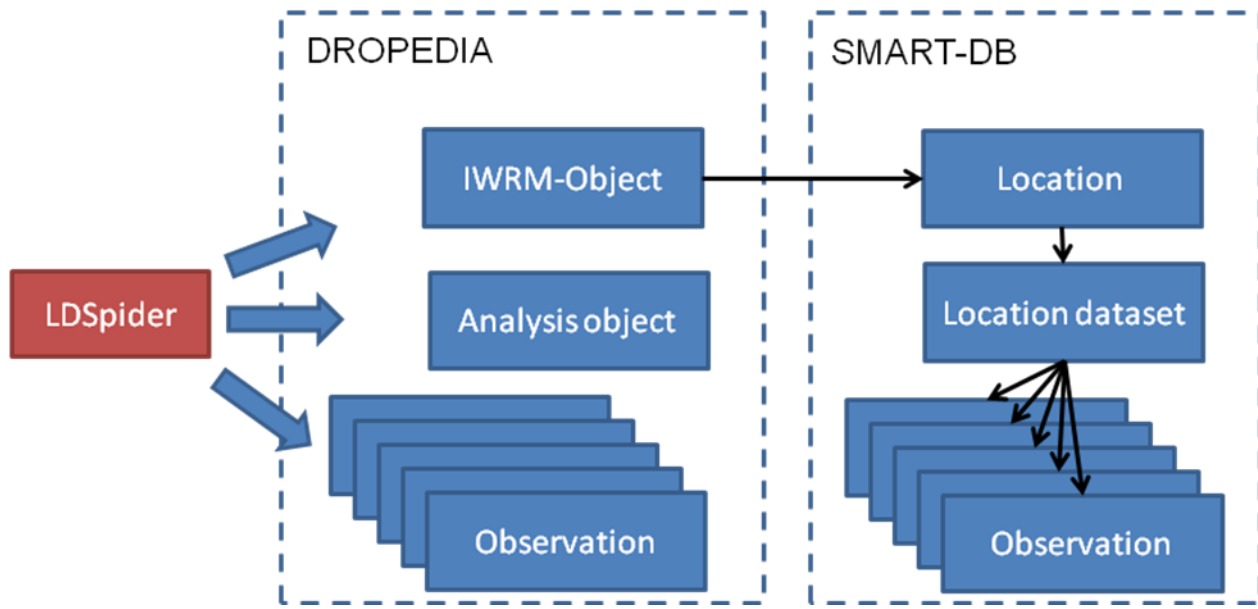


Figure 12 Illustration of crawling for data via LDSpider

An integrated knowledge base as a triple store has the following benefits:

- **Query interface:** The triple store provides a query interface Virtuoso SPARQL Query Editor¹¹ for structured queries using SPARQL (Working Group W. S., 2013). See Figure 13 for a screenshot of the query interface. For instance, the following query asks for all Catchments in the knowledge base:

Query for all catchments in the knowledge base:

```

prefix dropedia: <http://agkwebserver2.agk.uni-
karlsruhe.de/~dropedia/index.php/Special:URIResolver/>
select * where {?catchment a dropedia:Category-3ACatchment}

```

¹¹ <http://agkwebserver2.agk.uni-karlsruhe.de:8890/sparql>

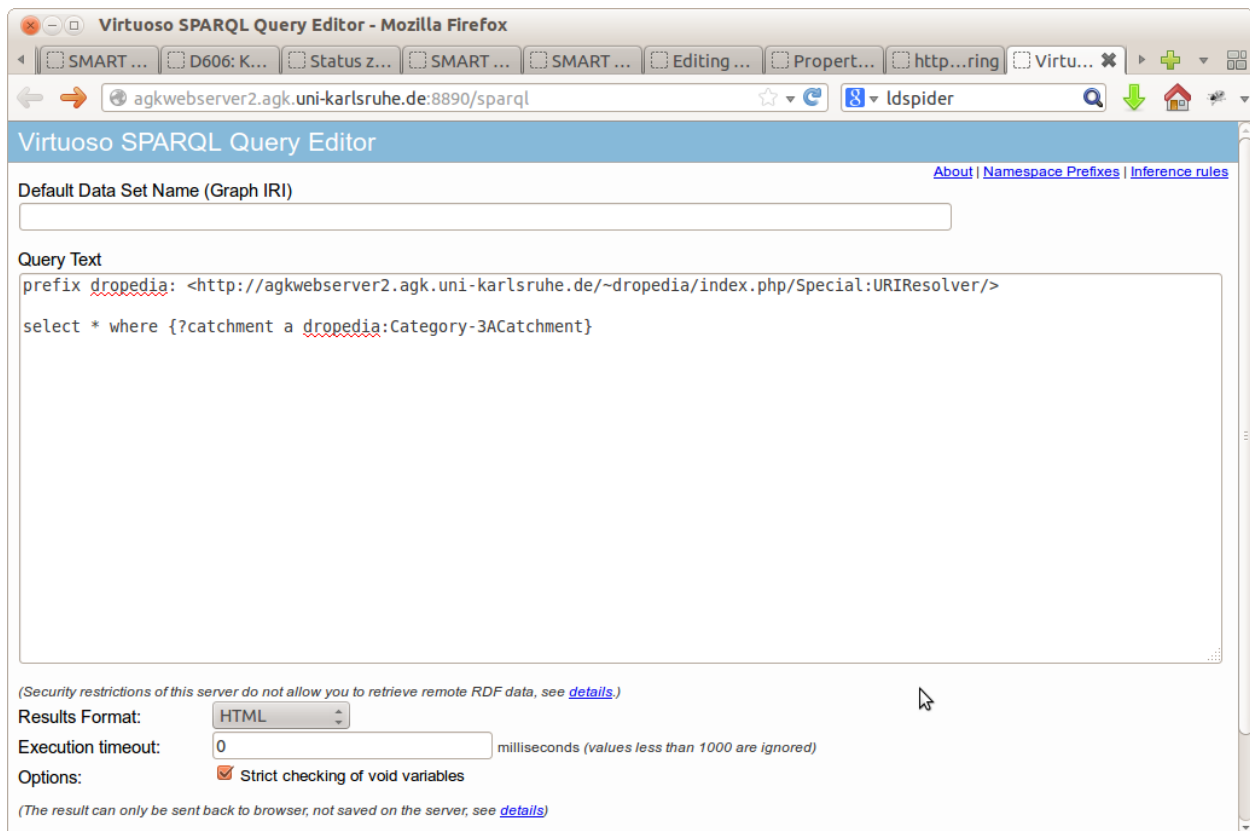


Figure 13 Screenshot of SPARQL Query Interface to SMART Knowledge Base

- OWL reasoning:** The triple store is able to interpret OWL axioms such as describing that Baqqouria_Spring in DROPEDIA is identical (`owl:sameAs`)¹² to AM0530 in SMART-DB. Therefore, the following two queries return the number (500) of all observations for the location, no matter what identifier for the location is used and no matter whether the observations were generated in DROPEDIA or SMART-DB.

First query using the identifier of DROPEDIA:

```
define input:same-as "yes"

prefix dropedia: <http://agkwebserver2.agk.uni-
karlsruhe.de/~dropedia/index.php/Special:URIResolver/>

prefix qb: <http://purl.org/linked-data/cube#>

prefix smart: <http://smartdbwrap.appspot.com/>

select count(?observation) where {?observation smart:location
dropedia:Baqqouria_Spring} limit 30
```

¹² Note, Open Virtuoso requires to state the `owl:sameAs` relationship between Baqqouria_Spring and AM0530 in both directions which we do using a SPARQL INSERT query.

Second query using the identifier of SMART-DB:

```
define input:same-as "yes"

prefix dropedia: <http://agkwebserver2.agk.uni-
karlsruhe.de/~dropedia/index.php/Special:URIResolver/>

prefix qb: <http://purl.org/linked-data/cube#>

prefix smart: <http://smartdbwrap.appspot.com/>

prefix smartdbwrap: <http://smartdbwrap.appspot.com/id/location/>

select count(?observation) where {?observation smart:location
smartdbwrap:AM0530} limit 30
```

The Virtuoso SPARQL Query Editor does not only allow returning the results as HTML or RDF, but also as CSV, a suitable format for import into WEAP. Yet, the query interface and SPARQL query language is more appropriate for technicians and requires information technology skills that scientists and decision makers typically do not possess. Rather, we provide this interface for programmers for innovative applications on top of the integrated SMART knowledge base. In the next section, we present the example of such an application, accessing the knowledge base.

1.5 SMART Data Explorer

In this section we present the SMART Data Explorer, an application on top of the SMART knowledge base for user-friendly, interactive data analysis and export.

The basic idea behind the SMART Data Explorer is to apply common OLAP and Data Warehousing concepts and technologies to provide scientists and decision makers an interface to analyse numeric data from the SMART knowledge base. For that, all observations from the IWRM ontology reusing the RDF Data Cube Vocabulary (Cyganiak & Reynolds, 2013) are modelled as one single Data Cube on which OLAP operations such as slice and dice can be issued (Kämpgen & Harth, 2011). See Figure 14 for an illustration of a Data Cube. Every observation from the SMART Knowledge Base is contained in a SMART-DB-DSD Data Cube¹³ that contains dimensions smart:location (e.g., Baqqouria Spring), smart:year, smart:month, smart:day (e.g., 1963-12-01), smart:analysis_Object (e.g., Q) and smart:analysis_Unit (e.g.,

¹³ <http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Special:URIResolver/SMART-2DDB-2DDSD>

m³/h) and two measures, AVG smart:obsValue and COUNT smart:obsValue (e.g., an average of 792 and a count of 498).




Figure 14 Illustration of a Data Cube

The architecture of the SMART knowledge base in Figure 11 already included components for the SMART Data Explorer. The triple store is accessed over its HTTP interface (Virtuoso SPARQL Query Editor), by a database driver (OLAP4LD). The database driver is accessed from a user-friendly and interactive user interface (Saiku) embedded in DROPIDIA. In the following, we describe in more detail those SMART Data Explorer components:

OLAP4LD

OLAP4LD¹⁴ is an Open-Source Java-based database driver that translates OLAP operations to SPARQL queries on a triple store (Kämpgen, O'Riain, & Harth, Interacting with Statistical Linked Data via OLAP Operations, 2012). For the SMART Data Explorer, we have connected OLAP4LD with the SMART Knowledge Base with the SMART-DB-DSD Data Cube.

Saiku

Saiku¹⁵ is an Open-Source OLAP client that allows user-friendly access to Data Cubes. Saiku can be embedded in wiki pages using the Widget Extension of MediaWiki. For the SMART Data Explorer, we have created a wiki page “SMART Data Explorer”¹⁶ in DROPIDIA.

In the following, we describe an example of using the SMART Data Explorer to export “Mean Discharge” data about Baqqouria Spring. The Baqqouria Spring wiki page on DROPIDIA¹⁷

¹⁴ <http://code.google.com/p/olap4ld/>

¹⁵ <http://analytical-labs.com/>

¹⁶ http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/SMART_Data_Explorer

contains SPARK queries on the SMART Knowledge Base 1) showing line diagram of the mean discharge time series and 2) showing the COUNT and AVERAGE of observation values for all analysis objects (indicators) in SMART-DB, see Figure 15 for the relevant parts of the DROPEDIA page.

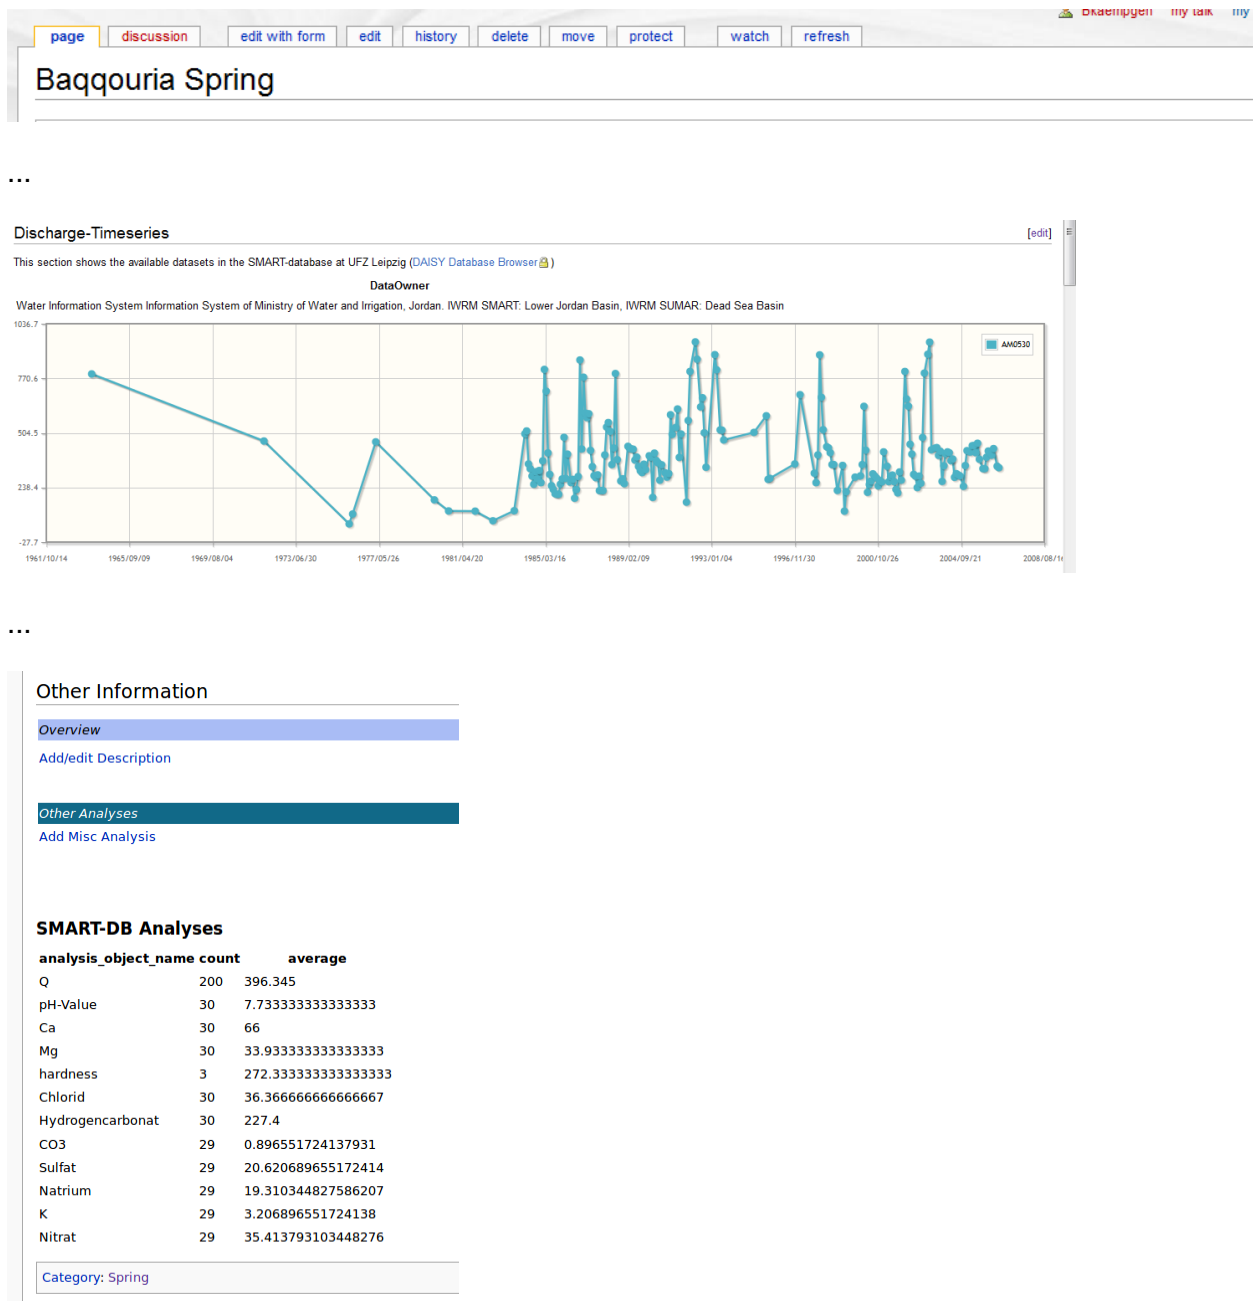


Figure 15 Snippets from DROPEDIA page of Baqqouria Spring with SMART-DB Queries

¹⁷ http://agkwebserver2.agk.uni-karlsruhe.de/~dropedia/index.php/Baqqouria_Spring

Also DROPEDIA can be used to collaboratively insert observations about Baqqouria Spring into the Knowledge Base, see Figure 16. Here, an Average Annual Discharge (identical to Mean Discharge and Q) for 2005 of 364 is documented.

The screenshot shows the DROPEDIA web interface in a Mozilla Firefox browser. The page title is 'Baqqouria Spring 10-year average discharge'. The interface includes a search bar, navigation links, and a table of monthly discharge measurements for 2005.

unit	jan	feb	mar	apr	may	jun	jul	aug	sep	oct	nov	dec	year
m ³ /h	295.7	268.8	322.6	349.5	457.0	483.9	470.4	403.2	349.5	376.3	322.6	268.8	364

Figure 16 Example of how to insert observations in DROPEDIA

Now, the SMART Data Explorer allows interactive and a user friendly exploration and exporting of such data. See Figure 17 for an example of using the SMART Data Explorer for querying the time series of Mean Discharge of Baqqouria Spring. Note the following:

- Every element mentioned in the SMART Data Explorer can be uniquely identified within the SMART knowledge base is available as Linked Data, i.e., can be browsed and documented. For instance, “refgovukyear:1963” is a resource provided as Linked Data by data.gov.uk, the British Data Portal. The full URI¹⁸ of refgovukyear:1963 provides additional information such the kind of calendar used.
- Since the SMART Data Explorer makes use of the reasoning capability of the triple store, identical elements are taken together. Therefore, only dropedia: Baqqouria_Spring and dropedia:Mean_Discharge are mentioned, although also smartdbwrap:AM0530 and

¹⁸ <http://reference.data.gov.uk/id/year/1963>

smartdbwrap:Q are queried for automatically. Consequently, the displayed result does not only contain the observations from SMART-DB, but also from DROPEdia.

- Data can be visualised in diagrams (naturally resembling the line chart on the DROPEdia page except for the observations integrated from DROPEdia), see Figure 18, and be exported to Microsoft Excel and CSV for import into WEAP.

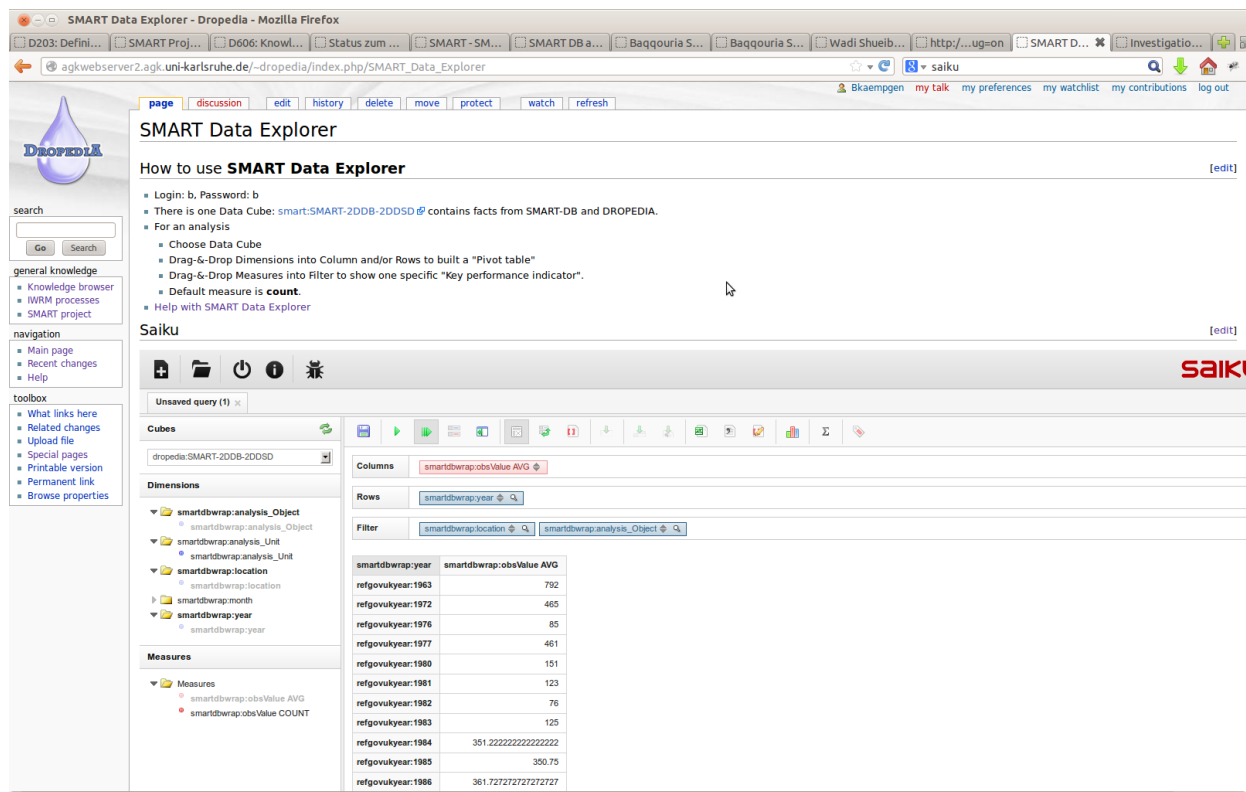


Figure 17 Screenshot of SMART Data Explorer with query for time series of Mean Discharge of Baqqouria Spring

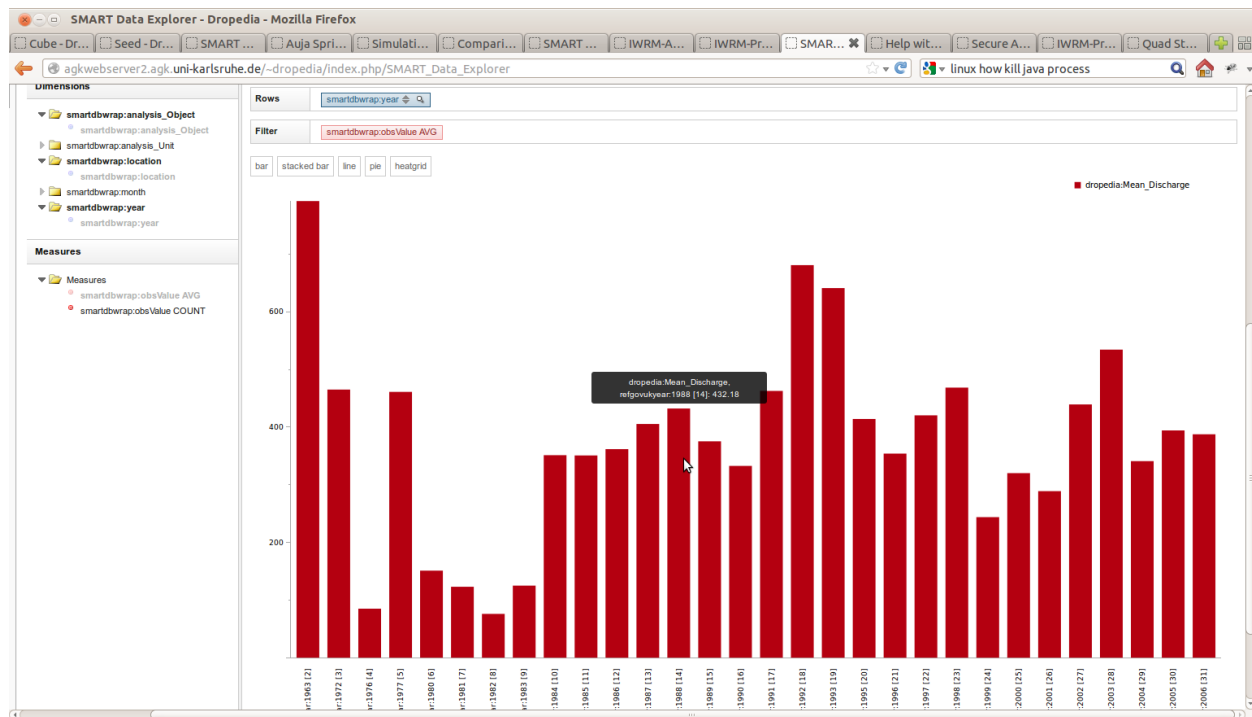


Figure 18 Bar chart of Baqqouria Spring Mean Discharge

The described concepts and technologies in this work have been presented to SMART project members in November 2012. Since then, only limited usage of the provided SMART Knowledge Base and SMART Data Explorer has been monitored. Workshops are now organised to help users. Therefore, a thorough user evaluation of the presented work is not available. Instead, in the following section, we describe a requirement coverage analysis for discussing our results.

4. Discussion of Results

In this section, we discuss our results via a requirement coverage analysis (see Table 4) comparing the derived requirements from our scenario with the capabilities of our results.

Table 4 Requirements Coverage Analysis

Requirement	Coverage
IWRM decision support elements	IWRM ontology describes processes, objectives, indicators and scenarios.
Links to unambiguously defined IWRM knowledge objects	IWRM ontology and DROPEDIA describe catchments (e.g., Wadi Shueib), municipalities (e.g., As-Salt) springs (e.g., Baqqouria Spring) and meteorological stations (e.g., Ira Meteo).
Decision makers and domain experts should be able to collaboratively work on IWRM decision support elements and IWRM knowledge objects.	Users can discuss, comment, re-use within DROPEDIA.
Domain experts should be able to collaboratively work on expert analyses.	Besides data-preprocessing analyses, estimations and assumption analyses, evaluations of indicators for scenarios, also, the actual decision matrix to be put in a multi-criteria-decision-analysis tool can be created and discussed in DROPEDIA.
All provided information should be published and accessible for the entire IWRM community.	The Linked Data principles allow uniquely linking to any resource for citations, feedback and related work.
DROPEDIA and SMART-DB are two separately created information sources that are to be integrated.	The SMART Knowledge Base stores data from both data sources in a data warehouse.

Also, WEAP information in the form of maps or numeric calculations to be integrated with DROPEdia or SMART-DB data.	WEAP KML files can be displayed in DROPEdia and links from WEAP objects to DROPEdia objects via URLs given. Analyses can be filled with observations.
Calculations (observations) needed for simulations in WEAP queried from DROPEdia or SMART-DB.	The SPARQL Query interface and the SMART Data Explorer allow querying integrated data from DROPEdia and SMART-DB for import into WEAP.
A DROPEdia webpage that describes an element also contained in SMART-DB shall give a tabular overview of the data that the SMART-DB stores about this element.	SMARTDBWRAP publishing SMART-DB as Linked Data and SPARK queries fulfil this requirement.
A DROPEdia webpage that describes an element also contained in SMART-DB shall give graphical insights of data available in the SMART-DB.	SMARTDBWRAP publishing SMART-DB as Linked Data and SPARK queries fulfil this requirement.
DROPEdia should allow to document and discuss datasets and give meta information such as the data owner, a time series visualisation, and an interactive visualisation.	An own category “Dataset” was created that allows to directly link to SMART-DB datasets for annotations (e.g., discussion, comments).
Not only humans should be able to reuse generated data, but also applications should be possible to develop that automatically process generated data and provide innovative analysis capabilities.	The SMART Data Explorer is an application showing the possibility and benefit of accessing SMART data in a machine-understandable way as Linked Data.

In summary, our results fulfil all requirements derived from our IWRM scenario.

5. Related Work

The German-Vietnamese water-related information system for the Mekong Delta (WISDOM) project provides a web-based information system (Gebhardt, et al., 2010). The system is based on PostgreSQL for geographical data management. Services are provided via representational state transfer (REST) and as such identify and allow access to resources similar to the Linked Data principles. However, the advantage of using REST services for efficient integration of WISDOM data sources with third-party data sources is not clear. There apparently has been research on ontology-based geospatial data management within WISDOM (Tran Thai, et al., 2010), however, it is not explained how this approach has contributed to the final information system. The graphical user interface promises metadata management, web sensor data analysis and thematic mapping. However, whereas all SMART IWRM tools are publicly available, the WISDOM information system is not. Also different from the WISDOM information system, the SMART Knowledge Base concept is focusing on data integration and making available data for third-party usage.

For the international project IWRM Indonesia, Bernhart et al. (Bernhart, Rösch, Vetter, & E., 2011) (Bernhart, Data storage of an IWRM project in a central Geographic Information System (GIS), 2011) have developed a system for data input and management. The system shall fulfil the following requirements: 1) central data storage to allow synergies and cost-efficiency, to avoid redundancies, without data owners losing control over their data, 2) easy visualisation of data and 3) usage of data for GIS-teaching activities. Data heterogeneity is considered one of the most difficult challenges for GIS data management. The presented system provides web-based forms for data input. Forms can be extended and modified by a system administrator. With respect to data-storage and data-input, the IWRM Indonesia concept is similar to the SMART concept. SMART, however, also focuses on 1) data integration of different data sources such as SMART-DB with sensor data and DROPEDIA with semi-structured data, 2) decision support (using multi-criteria-decision analysis) and data mining (using OLAP) as well as 3) making available IWRM data via a machine-understandable ontology on the Semantic Web for third-party consumption and applications.

6. Conclusions

Integrated Water Resources Management (IWRM) has great potential to improve water scarcity situations in regions such as the Lower Jordan Valley, the study area of the BMBF SMART project. However, IWRM is 1) lacking concrete operational guidelines to help scientists contributing to IWRM processes and 2) missing knowledge management methods and tools to share and integrate information from social, economical and ecological sciences. In this work, we try to overcome those problems via semantic technologies. Based on IT-well-known Linked Data Principles, we have designed and developed an integrated SMART Knowledge Base with data from different data sources and defined data exchange protocols to keep the data up-to-date, to integrate the data with other data sources, and to allow efficient and user-friendly access. More concretely, in this work, we have presented 1) an IWRM ontology about IWRM knowledge and decision support, 2) mechanisms to publish the SMART platform DROPIEDIA and SMART-DB as Linked Data reusing the IWRM ontology, 3) an ETL pipeline to create and maintain an IWRM knowledge base with integrated SMART data and 4) the SMART Data Explorer to analyse data from the IWRM knowledge base.

We have evaluated our work with a coverage analysis of technical requirements derived from a real-world IWRM decision process for the Wadi Shueib catchment in Jordan. Although our concept fulfils the technical requirements, for a thorough user study more systematic training of potential users would be required.

We see potential in our contributions to data exchange formats and data integration regarding the ongoing “Open Data” trend. For instance, more and more Public Sector Information is made available, e.g., from data.gov.uk (UK), data.gov (US) and govdata.de (DE) and promises useful to be integrated with IWRM related data sources. Also, more and more institutions such as DataCite¹⁹ and Pangea²⁰ help scientists to not only publish their analysis results but also the collected and analysed raw (or also pre-processed) data available for citations, reproduction and further analysis. In this work we have presented concepts and technologies allowing the efficient integration and analysis of such growing number of distributed data sources.

¹⁹ <http://www.datacite.org/>

²⁰ <http://pangaea.de/>

7. Bibliography

Bensabat, J., & Dror, S. (2013). Multi-Criteria Optimization (MCO) Module - Version 1.2. SMART.

Bensabat, J., Kitron-Belinkov, M., Milgrom, T., & Dror, S. (2012). DSS Development - SMART Deliverable 601.

Berners-Lee, T. (27. 07 2006). Linked Data. Abgerufen am 06. 04 2013 von <http://www.w3.org/DesignIssues/LinkedData.html>

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. Scientific American .

Bernhart, F. (2011). Data storage of an IWRM project in a central Geographic Information System (GIS). International Conference on Integrated Water Resources Management. Dresden.

Bernhart, F., Rösch, N., Vetter, M., & E., R. (2011). Datenorganisation eines interdisziplinären Verbundprojektes. Tagungsband des Entwicklerforums: Geodäsie und Geoinformationstechnik. Berlin: Shaker Verlag Aachen.

Cyganiak, R., & Reynolds, D. (2013). The RDF Data Cube Vocabulary. W3C.

Gebhardt, S., Wehrmann, T., Klinger, V., Schettler, I., Huth, J., Künzer, C., et al. (2010). Improving data management and dissemination in web based information systems by semantic enrichment of descriptive data aspects. Computers & Geosciences .

Geyer, S., & Siebert, C. (2009). Data and Information System (DAISY).

Hitzler, P., Krötzsch, M., & Rudolph, S. (2010). Foundations of Semantic Web Technologies. CRC Press.

Kämpgen, B., & Harth, A. (2011). Transforming Statistical Linked Data for Use in OLAP Systems. I-SEMANTICS 2011, 7th Int. Conf. on Semantic Systems. ACM.

Kämpgen, B., O'Riain, S., & Harth, A. (2012). Interacting with Statistical Linked Data via OLAP Operations. Proceedings of Interacting with Linked Data (ILD 2012), workshop co-located with the 9th Extended Semantic Web Conference (S. 36-49). CEUR-WS.

Manola, F., & Miller, E. (2004). RDF Primer. W3C.

Riepl, D. (2013). Knowledge-Based Decision Support for Integrated Water Resources Management with an application for Wadi Shueib, Jordan. Karlsruhe.

Riepl, D., & Kämpgen, B. (2013). Knowledge Management System - Deliverable 606. SMART.

Riepl, D., Wolf, L., Kaempgen, B., & Heinz, H. (2010). Making informed decisions – a collaborative and knowledge based IWRM planning exercise in WadiShueib, Jordan. Proceedings IWRM 2010 . Karlsruhe.

Tran Thai, B., Wehrmann, T., Gebhardt, S., Klinger, V., Huth, J., Vo Quoc, T., et al. (2010). Proceedings of the 31st Asian Remote Sensing Conference.

Völkel, M., Krötzsch, M. V., Haller, H., & Studer, R. (2006). Semantic Wikipedia. WWW 2006.

Working Group, W. O. (2012). OWL 2 Web Ontology Language - Document Overview (Second Edition). W3C.

Working Group, W. S. (2013). SPARQL 1.1 Overview. W3C.